Amendments to the Specification

Replace the paragraph beginning on page 2 line 9 with the following:

The commercial viability of a storage system reflects the architectural decisions and component selections made by the designer to provide a desired level of fault tolerance, storage capacity, operating life, and data availability. ~~Components with very long MTBF (mean time between failure) ratings may adversely affect system cost.~~

Replace the paragraph beginning on page 2 line 17 with the following:

Embodiments of the present invention ~~furnishes~~ furnish redundant storage system architectures and isolation methods that provide fault tolerance in data storage systems and that can be employed to eliminate single points of failure.

Replace the paragraph beginning on page 5 line 6 with the following:

FIGURE 8 depicts a loop bypass storage system with ~~two~~ dual ported drives connected to each port of a port bypass controller.

Replace the paragraphs beginning on page 8 line 1 with the following:

FIGURE 3 depicts a loop storage system architecture. System **300** comprises host **302**, disc array controller **304**, bus **306**, and drive array **308**. Disc array controller **304** is connected to host **302** by one or more buses. Bus 306 serially interconnects disc array controller 304 and each of the drives of drive array **308** in a loop. Disc array controller **304** and each drive of drive array **308** have ~~on~~ an input port and an output port ~~to~~ connected to form the loop of bus **306**. The system of figure 3 can continue to operate if a disc failure

occurs that does not affect bus operation. The failure of the bus, controller, or a disc failure that interrupts bus operation results in loss of data availability, requiring repair of the bus, controller, or disc drive, or installation of drives in another fixture to access data.

Replace the paragraphs beginning on page 8 line 11 with the following:

FIGURE 4 depicts a storage system architecture employing switched single-ported disc drives. System 400 comprises host 402, disc controller "A" 404, disc controller "B" 406, switch control 408, bus "A" 410, bus "B" 412, disc drives 414-422 and switching devices 424-432. Disc controller "A" 404 and disc controller "B" 406 are connected to host 402 by one or more buses and are dual ported ~~that~~ so that they each provide two disc drive buses. Bus "A" 410 and bus "B" 412 are connected to both disc controller "A" 404 and disc controller "B" 406. In an alternative embodiment (not depicted), two single port disc controllers can be used wherein a first disc controller provides communication on bus "A" 410 and a second disc controller provides communications on bus "B" 412. Switching devices 424-432 are controlled by switch control 408 and independently connect drives 414-422 to bus "A" 410 or bus "B" 412. Switching devices 424-432 may be any type of switching devices including but not limited to cross-point switches, port multiplexers and the like. Switch control may comprise one or more buses that connect switching devices 424-432 to host 402 and may comprise an I2C bus, RS232, or any other serial or parallel buses. Alternatively, switching devices may be controlled by disc controller "A" 404, disc controller "B" 406, or both. In another embodiment, switch control may employ bus "A" 410 and/or bus "B" 412. As such, switching devices may be controlled directly by host 402, by host 402 through disc controller "A" ~~410~~ 404 or disc controller "B" ~~412~~ 406, or may be controlled by

disc controller "A" ~~410~~ 404 or disc controller "B" ~~412~~ 406. The architecture of figure 4 may employ a larger number of discs and switching devices than depicted. Switching devices can be individually configured for each drive such that each drive employs either bus "A" 410 or bus "B" 412. This allows communication to be maintained in the event of a bus failure, and allows loads to be balanced between buses. The architecture of figure 4 provides continued operation in the event of a bus, disc, or controller failure. Switching devices 424-432 may also allow disc drives to be isolated from both buses. In the event of a disc failure, or a disc failure that corrupts bus operation, an associated switching device may be configured to disconnect the drive from both buses. The switching methods shown in figure 4 may be applied to dual ported drives where each port of each drive may be selectively connected to bus "A" 410, bus "B" 412, or may be disconnected from both buses. Alternatively, a third bus may be employed to provide higher transfer rates in the event of a bus failure.

Replace the paragraph beginning on page 10 line 12 with the following:

The architecture of figure 5 allows system operation to continue after the failure of one or more disc controllers, disc drives, or buses. Additionally, the architecture of figure 5 allows data loads to be distributed among disc controllers and buses to optimize performance. Depending upon the number of disc drives, and the data rates of disc drives, the buses, and, disc controllers, the architecture of figure 5 may provide near optimum performance following the failure of a disc drive, bus, or disc controller. As such the above architecture may be employed in systems where continued high performance is desired following failure of a bus ~~of~~ or disc controller.

Replace the paragraph beginning on page 11 line 4 with the following:

Loop bypass methods may be employed to isolate one or more drives. More than one drive may be connected to each port of a port bypass controller. Figure 7 depicts a loop bypass storage system with two drives connected to each bypass controller port. System 700 comprises host 702, disc controller 704, disc drives 706-724, port bypass controller 726, and bus 728. Drives are arranged in pairs such that drives 706,708 are connected to a first port of port bypass controller 726, drives 710,712 are connected to a second port, drives 714-716, are connected to another port, drives 718,720 are connected to yet another port, and drives 722,724 are connected to still another port. Bus 728 connects disc controller 704 to port bypass controller 726. In an alternative embodiment, two buses may connect the disc controller and port bypass controller, providing redundancy in the event of a bus failure. Any ~~or~~ of the ports of port bypass controller 726 may be configured to allow signals to pass through the two drives connected to the port or to bypass the port, providing isolation in the event of a drive failure, or drive failure that corrupts the bus. While figure 7 depicts two drives connected to each port of port bypass controller 726, more than two drives may be connected within the scope of the present invention. While figure 7 employs a port bypass controller, any devices and configuration thereof that produce the described function may be employed.

Replace the paragraphs beginning on page 11 line 22 with the following:

Loop bypass architectures may employ a plurality of drives connected to each port wherein each drive is dual ported. Figure 8 depicts a loop bypass storage system with ~~two~~ dual ported drives connected to each port bypass controller 808, 812. System 800 comprises

host 802, disc controller 804, disc controller 806, port bypass controller 808, bus 810, port bypass controller 812, bus 814 and disc drives 816-824. Disc controller 804 and disc controller 806 are each connected to host 802 by one or more buses. Disc controller 804 is connected to port bypass controller 808 through bus 810. Disc controller 806 is connected to port bypass controller 812 through bus ~~b~~ 814. In an alternative embodiment, more than one bus may connect disc controller 804 to port bypass controller 808, and more than one bus may connect disc controller 806 to port bypass controller 812. In another embodiment, each disc controller may connect to both port bypass controllers. Disc drives 816-~~814~~824 are dual ported and each drive has a first port connected to port bypass controller 808 and a second port connected to port bypass controller 812. As such, each disc drive may be individually configured to connect to a loop formed by bus 810 on one port, or bus 814 on the second port of the drive, or both buses. In the event of a drive failure, or drive failure that corrupts bus signals, the drive may be isolated through configuration of port bypass controller 808 or port bypass controller 812, or configuration of both port bypass controllers. In the event of a disc controller failure, bus failure, connector failure, or port bypass controller failure, data from drives may be accessed using the functioning disc controller, bus, or port bypass controller.

Replace the paragraphs beginning on page 12 line 11 with the following:

Two or more dual ported disc drives may be connected to each port of a port bypass controller. Figure 9 depicts a loop bypass storage system with two dual ported drives connected to each ~~port of a~~ port bypass controller 908, 934. System 900 comprises host 902, disc controller 904, bus 906, port bypass controller 908, disc drives 910-928, disc controller 930, bus 932, and port bypass controller 934. Disc controller 904 and disc controller 930 are

connected to host 902 by one or more buses. Disc controller 904 is connected to port bypass controller 908 through bus 906. Disc controller 930 is connected to port bypass controller 934 through bus 932. Disc drives 910-928 are dual ported and each drive has a first port connected to port bypass controller 908 and a second port connected to port bypass controller 934. In an alternative embodiment, disc controller 904 is also connected to port bypass controller 934 and disc controller 930 is also connected to port bypass controller 908. Port bypass controllers 908 and 934 are individually configurable to provide a connection to a disc drive port or to bypass a connection to a disc drive, allowing each disc drive to be isolated in the event or a drive failure or a failure that corrupts the port connection. Since disc drives are dual ported and two port bypass controllers are employed, the system of figure 9 provides continued operation in the event of a disc controller failure, bus failure, or disc drive failure.

Replace the paragraphs beginning on page 13 line 15 with the following:

Figure 11 depicts another multi-path redundant storage system. System ~~1000~~ 1100 comprises system interface 1102, system bus "A" 1104, system bus "B" 1106, interface controller "A" 1108, interface controller "B" 1110, interface bus "A" 1112, interface bus "B" 1114, disc controller "A" 1116, disc controller "B" 1118, fabric bus "A" 1120, fabric bus "B" 1122, fabric "A" 1124, fabric "B" 1126, fabric control bus "A" 1128, fabric control bus "B" 1130, and drive groups 1132-1140. Interface controller "A" 1108 and interface controller "B" 1110 connect to a system through system bus "A" 1104 and system bus "B" 1106. The two system buses provide redundant communication paths, allowing continued communication with both interface controllers in the ~~vent~~ event that one of the system buses

fails. Interface controller "A" 1108 and interface controller "B" 1110 connect to disc controller "A" 1116 and disc controller "B" 1118 through interface bus "A" 1112 and interface bus "B" 1114 that allow continued communication between either interface controller and either disc controller in the event that one of the interface buses fails. Disc controller "A" 1116 and disc controller "B" 1118 are connected to fabric "A" 1124 and fabric "B" 1126 through fabric bus "A" 1120 and fabric bus "B" 1122, providing continued communication between either disc controller and either fabric in the event that one of the fabric buses fails. Fabric control bus "A" 1128 and fabric control bus "B" 1130 provide redundant control paths from interface controller "A" 1108 and interface controller "B" 1110 to fabric "A" 1124 and fabric "B" 1126 and allow configuration of either fabric by either interface controller in the event that either fabric control bus fails. Fabric "A" 1124 is connected to each drive group of drive groups 1132-1140 by separate connection. A drive group comprises one or more drives connected to a fabric by one connection. Drives in the drive groups are dual ported. Fabric "B" 1126 is connected to each drive group of groups 1132-1140 by separate connection. Fabric "A" 1124 connects to one port of the dual ported drive or drives comprising each drive group and fabric "B" 1126 connects to a second port of the dual ported drive or drives comprising each group. The duality of system buses, interface buses, fabric buses, fabric control buses, and drive group connections provides isolation or a redundant path for every data path in the system. The duality of interface controllers, disc controllers, and fabrics, in conjunction with the duality of buses, provides continued operation in the event of a failure of an interface controller, disc controller, or fabric. As such the system depicted in figure 11 has no single point of failure relative to buses, controllers, or fabrics.

Replace the paragraphs beginning on page 14 line 18 with the following:

In addition to buses, connectors, disc drives, fabrics and controllers, isolation and redundancy methods may be further applied to power distribution in a storage system such that the system has no single point of failure that might render the system inoperative. Figure 12 depicts multi-path redundant storage system power distribution. Power is supplied to the system through connector 1202. Alternatively, more than one connector may be employed. More than one contact pin within a connector may provide a like voltage, providing a duality of paths in the event that one pin fails to make connection or has higher than desired resistance. Power bus "A" 1204 provides power to local regulator 1208, local regulator 1212, and optionally may provide power to one or more additional local regulators as indicated by local regulator 1216. Local regulator 1208 provides power to fabric "A" 1206. Local regulator 1212 provides power to fabric "B" 1210. Optional regulator 1216 may provide power to disc controller 1214. Other local regulator (not depicted) may provide power to additional disc controllers and to interface controllers, discrete circuitry, or other circuitry such as environmental monitors, for example. Local regulators may be employed to provide power regulated to a desired voltage to components such as integrated circuits that consume relatively low power as compared to disc drives. Systems having redundant interface controllers, disc controllers, and fabrics may employ local regulators for each component, providing continued system operation in the event that a single regulator fails since the redundant component may be employed to access data. Connector 1202 of figure 12 also provides one or more pins connected to power bus "B" 1218. Power bus "B" 1218 provides power to voltage regulators 1220 and 1222. Regulators 1220 and 1222 are connected in a manner that allows power to be provided by either regulator and may include isolation

circuitry such as diodes or other components. Alternatively, regulators 1220 and 1222 may include input signals that may enable or disable each regulator. Regulators may be controlled by writeable registers, I2C buses, or other signal lines. Voltage regulators 1220 and 1222 provide regulated power to control 1224, control 1228, and optionally to one or more additional controls as indicated by control 1232. Control 1224 controls power to disc group 1226. Control 1228 controls power to disc group 1230. Control 1232 provides power to disc group 1234. Additional control units (not depicted) may control power to additional disc groups, or to other components such as environmental monitors, fans, or other components. Controls 1224, 1228, 1232 and other controls may comprise switches, fuses, breakers, transistors (including field effect transistors, SCRs (silicon controlled rectifiers) or any other devices employed to selectively apply power to a disc group or other components. Controls may include current and/or voltage sensing and may operate in an automatic manner or in response to a control signal. Figure 12 illustrates that methods of power redundancy and isolation may be applied to data storage system components such that data remains available following the failure of a regulator, and that power to one or more disc drives in a group containing a failed drive may be shut off to conserve power in the system or to isolate components drawing excessive power. As previously noted, data from a failed drive or drive group may be copied or reconstructed and saved using spare capacity of functioning drives. As such, embodiments of the present invention can provide a data storage system that has no single point of failure that would result in data loss.

Replace the paragraph beginning on page 17 line 12 with the following:

Figure 13 depicts steps performed by system configuration computer program code operating in a host and/or disc controller. The process of figure 13 is applicable to systems like that shown in figures 10 and/or 11. Process 1300 begins at step 1302 where a check is performed to determine if an error condition exists. An error condition may comprise an error such as a read or write error, for example, detected by a disc drive, disc controller, or host system. If the error is detected by a disc drive, the error may be reported to a disc controller and may be checked by a disc controller and/or may be forwarded to a host system. If a disc controller detects an error, the error may be checked and/or may be forwarded to a host system. Alternatively, an error may be detected by a host system. At step 1304, a test may be performed to determine if the host can communicate with interface controller "A" using system bus "A". At step 1306, a test may be performed to determine if the host can communicate with interface controller "A" using system bus "B". At step 1308, a test may be performed to determine if the host can communicate with interface controller "B" using system bus "A". At step 1310, a test may be performed to determine if the host can communicate with interface controller "B" using system bus "B". Steps 1304-1310 determine if a host or other system is able to communicate with interface controller "A and interface controller "B" using both system bus "A" and system bus "B". At step 1312, any errors detected in steps 1304-1310 are reported to a host or other system. At step 1314, a check is performed, such as reviewing reported errors, for example, to determine if the host or other system is able to communicate with at least one interface controller. If the host or other system is not able to communicate with at least one interface controller, the process ends at step 1316. If the check performed at step 1314 determines that the host or other system is able to communicate with at least one interface controller, the process continues at step 1318

where a test is performed to determine if disc controller "A" can be accessed using interface bus "A". This test may comprise reading disc controller registers. At step 1320, a test is performed to determine if disc controller "A" can be accessed using interface bus "B". At step 1322, a test is performed to determine if disc controller "B" can be accessed using interface bus "A". At step 1324, a test is performed to determine if disc controller "B" can be accessed using interface bus "B". At step 1326, any errors detected in steps 1318-1324 are reported. At step 1326, test results are checked to determine if at least one disc controller can be accessed. If no disc controllers can be accessed, the process ends at step 1330. If at least one disc controller can be accessed, the process continues at step 1332 where a test is performed to determine if fabric "A" can be accessed using fabric bus "A". At step 1334 a test is performed to determine if fabric "A" can be accessed using fabric bus "B". At step 1336 a test is performed to determine if fabric "B" can be accessed using fabric bus "A". At step 1338 a test is performed to determine if fabric "B" can be accessed using fabric bus "B". At step 1340, any errors detected in steps 1332-1338 are reported. At step 1342, test results are check to determine if at least one fabric is accessible. If no fabrics are accessible, the process ends at step 1344. If at least one fabric is accessible, the process continues at step 1346. At step 1346 a test is performed to determine if fabric "A" can access all attached drives. Such tests may comprise reading and/or writing drive registers and/or reading and/or writing data to the drive media. If not all drives are accessible or are not operating properly, fabric "A" may be configured to isolate one or more drives in step 1348 and then the process continues at step 1350. If the test performed in step 1346 determines all drives are accessible and are operating properly, the process continues at step 1350. At step 1350, a test is performed to determine if fabric "B" can access all attached drives. If some drives are not

accessible, or are not operating properly, fabric "B" may be configured to isolate one or more drives in step 1352 and the process then continues at step 1354. At step 1354, data from inaccessible or failed drives may be reconstructed or copied and stored on other drives or may be stored on another system such that fault tolerance is provided.[[.]] I/O commands may be remapped to utilize functioning interface controllers, disc controllers, or fabrics, as identified by ~~pervious~~ previous tests. The process then ends at step 1356. If the test performed in step 1350 determines that all drives are accessible and operating properly, the process ends at step 1356. The results of tests performed may also be employed to configure power circuitry such as depicted in figure 12 such that power is not applied to failed components. The tests performed, the order of tests performed, configuration of fabrics and reconstruction of data and remapping of I/Os may be varied depending on the architecture of the storage system including the number of host buses, interface controllers, disc controllers, number and type of fabrics, and number of disc drives including the number of disc drives attached to each port of the fabric or fabrics. The type of error reported may be used to select a test or set of tests. Alternatively, following a reported error, a range of tests may be run to determine the overall condition of a storage subsystem. A hierarchical order of tests may exist wherein operation of various system components is performed in a predetermined order. The tests performed in figure 13 may be executed by a host or other system, or may be executed by components within a storage subsystem. Computer program code performing tests may be resident in individual components of the system or may be transferred from other systems or other components. Tests may include execution of self-test computer program code in components. For example, disc drives may include a power-on self test

routine and such ~~routing~~ routine may be invoked as part of the tests performed in figure 13 to check operation of disc drives.